

 **PORTAL**  
USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login  
 Search:  The ACM Digital Library  The Guide  
 +InfiniBand, +database replicate, update

The ACM Digital Library | Log In

Feedback Report a problem Satisfaction survey

Terms used InfiniBand database replicate update

Found 18 of 171,143

Sort results by relevance   
 Display results expanded form    
 Open results in a new window

Try an Advanced Search  
 Try this search in The ACM Guide

Results 1 - 18 of 18

Relevance scale 

### 1 Cluster communication protocols for parallel-programming systems

 Kees Verstoep, Raoul A. F. Bhoedjang, Tim Rühl, Henri E. Bal, Rutger F. H. Hofman  
 August 2004 **ACM Transactions on Computer Systems (TOCS)**, Volume 22 Issue 3

Publisher: ACM Press

Full text available:  pdf(1.29 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Clusters of workstations are a popular platform for high-performance computing. For many parallel applications, efficient use of a fast interconnection network is essential for good performance. Several modern System Area Networks include programmable network interfaces that can be tailored to perform protocol tasks that otherwise would need to be done by the host processors. Finding the right trade-off between protocol processing at the host and the network interface is difficult in general. In ...

**Keywords:** Clusters, parallel-programming systems, system area networks

### 2 Empirical evaluation of multi-level buffer cache collaboration for storage systems

 Zhipeng Chen, Yan Zhang, Yuanyuan Zhou, Heidi Scott, Berni Schiefer  
 June 2005 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2005 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '05**, Volume 33 Issue 1

Publisher: ACM Press

Full text available:  pdf(379.25 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

To bridge the increasing processor-disk performance gap, buffer caches are used in both storage clients (e.g. database systems) and storage servers to reduce the number of slow disk accesses. These buffer caches need to be managed effectively to deliver the performance commensurate to the aggregate buffer cache size. To address this problem, two paradigms have been proposed recently to *collaboratively* manage these buffer caches together: the ***hierarchy-aware caching*** maintains ...

**Keywords:** collaborative caching, database, file system, storage system

### 3 WireGL: a scalable graphics system for clusters

 Greg Humphreys, Matthew Eldridge, Ian Buck, Gordan Stoll, Matthew Everett, Pat Hanrahan  
 August 2001 **Proceedings of the 28th annual conference on Computer graphics and interactive techniques**

**Publisher:** ACM Press

Full text available:  pdf(333.39 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We describe WireGL, a system for scalable interactive rendering on a cluster of workstations. WireGL provides the familiar OpenGL API to each node in a cluster, virtualizing multiple graphics accelerators into a sort-first parallel renderer with a parallel interface. We also describe techniques for reassembling an output image from a set of tiles distributed over a cluster. Using flexible display management, WireGL can drive a variety of output devices, from standalone displays to tiled displ ...

**Keywords:** cluster rendering, parallel rendering, remote graphics, scalable rendering, tiled displays, virtual graphics

**4 Power management and voltage scaling: Power-aware code scheduling for clusters of active disks** 



S. W. Son, G. Chen, M. Kandemir

August 2005 **Proceedings of the 2005 international symposium on Low power electronics and design ISLPED '05**

**Publisher:** ACM Press

Full text available:  pdf(287.34 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper, we take the idea of application-level processing on disks to one level further, and focus on an architecture, called Cluster of Active Disks (CAD), where the storage system contains a network of parallel "active disks." Each individual active disk (which includes an embedded processor, disk(s), caches, memory, and interconnect) can perform some application level processing; but, more importantly, the active disks can collectively perform parallel Input/Output (I/O) and processing, ...

**Keywords:** cluster of active disks (CAD), compiler, scheduling

**5 Promises and reality: Performance measurements of a user-space DAFS server with a database workload** 



Samuel A. Fineberg, Don Wilson

August 2003 **Proceedings of the ACM SIGCOMM workshop on Network-I/O convergence: experience, lessons, implications**

**Publisher:** ACM Press

Full text available:  pdf(366.48 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We evaluate the performance of a user-space Direct Access File System (DAFS) server and Oracle Disk Manager (ODM) client using two synthetic test codes as well as the Oracle database. Tests were run on 4-processor Intel Xeon-based systems running Windows 2000. The systems were connected with ServerNet II, a Virtual Interface Architecture (VIA) compliant system area network. We compare the performance of DAFS/ODM and local-disk based I/O, measuring I/O bandwidth and latency. We also compare the r ...

**Keywords:** DAFS, Database, File Systems, I/O, Networks, Performance Evaluation, RDMA

**6 High Resolution Aerospace Applications using the NASA Columbia Supercomputer** 

Dimitri J. Mavriplis, Michael J. Aftosmis, Marsha Berger

November 2005 **Proceedings of the 2005 ACM/IEEE conference on Supercomputing SC '05**

**Publisher:** IEEE Computer Society

Full text available:  [pdf\(3.10 MB\)](#) Additional Information: [full citation](#), [abstract](#)  
[Publisher Site](#)

This paper focuses on the parallel performance of two high-performance aerodynamic simulation packages on the newly installed NASA Columbia supercomputer. These packages include both a high-fidelity, unstructured, Reynolds-averaged Navier-Stokes solver, and a fully-automated inviscid flow package for cut-cell Cartesian grids. The complementary combination of these two simulation codes enables high-fidelity characterization of aerospace vehicle design performance over the entire flight envelope t ...

**Keywords:** NASA Columbia, SGI Altix, scalability, hybrid programming, unstructured, computational fluid dynamics, OpenMP

## 7 Experiences with VI communication for database storage

Yuanyuan Zhou, Angelos Bilas, Suresh Jagannathan, Cezary Dubnicki, James F. Philbin, Kai Li  
 May 2002 **ACM SIGARCH Computer Architecture News , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02 , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02**, Volume 30 Issue 2

**Publisher:** IEEE Computer Society, ACM Press

Full text available:  [pdf\(1.29 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)  
[Publisher Site](#)

This paper examines how VI-based interconnects can be used to improve I/O path performance between a database server and the storage subsystem. We design and implement a software layer, DSA, that is layered between the application and VI. DSA takes advantage of specific VI features and deals with many of its shortcomings. We provide and evaluate one kernel-level and two user-level implementations of DSA. These implementations trade transparency and generality for performance at different degrees ...

**Keywords:** Storage system, cluster-based storage, Database storage, storage area network, User-level Communication, Virtual Interface Architecture, processor overhead

## 8 Session 9: operating systems: High performance support of parallel virtual file system (PVFS2) over Quadrics

Weikuan Yu, Shuang Liang, Dhabeleswar K. Panda  
 June 2005 **Proceedings of the 19th annual international conference on Supercomputing ICS '05**

**Publisher:** ACM Press

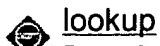
Full text available:  [pdf\(256.45 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

Parallel I/O needs to keep pace with the demand of high performance computing applications on systems with ever-increasing speed. Exploiting high-end interconnect technologies to reduce the network access cost and scale the aggregated bandwidth is one of the ways to increase the performance of storage systems. In this paper, we explore the challenges of supporting parallel file system with modern features of Quadrics, including user-level communication and RDMA operations. We design and implemen ...

**Keywords:** RDMA, parallel IO, parallel file system, quadrics, zero-copy

## 9

### I/O: miNI: reducing network interface memory requirements with dynamic handle



Reza Azimi, Angelos Bilas

June 2003 **Proceedings of the 17th annual international conference on Supercomputing**

**Publisher:** ACM Press

Full text available: [pdf\(289.75 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Recent work in low-latency, high-bandwidth communication systems has resulted in building user--level Network Interface Controllers (NICs) and communication abstractions that support direct access from the NIC to applications virtual memory to avoid both data copies and operating system intervention. Such mechanisms require the ability to directly manipulate user--level communication buffers for delivering data and achieving protection. To provide such abilities, NICs must maintain appropriate t ...

**Keywords:** parallel architectures, system area networks

**10 Editors' choice awards 2005**

Don Marti

August 2005 **Linux Journal**, Volume 2005 Issue 136

**Publisher:** Specialized Systems Consultants, Inc.

Full text available: [html\(21.56 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

We want our servers stable, our graphics non-jagged and our drivers GPL. Here's a shopping-cart load of the stuff that makes us happy.

**11 upFront**

Linux Journal Staff

July 2004 **Linux Journal**, Volume 2004 Issue 123

**Publisher:** Specialized Systems Consultants, Inc.

Full text available: [html\(10.72 KB\)](#) Additional Information: [full citation](#)

**12 Runtime Compression of MPI Messages to Improve the Performance and Scalability of Parallel Applications**

Jian Ke, Martin Burtscher, Evan Speight

November 2004 **Proceedings of the 2004 ACM/IEEE conference on Supercomputing**

**Publisher:** IEEE Computer Society

Full text available: [pdf\(111.21 KB\)](#) Additional Information: [full citation](#), [abstract](#)

Communication-intensive parallel applications spend a significant amount of their total execution time exchanging data between processes, which leads to poor performance in many cases. In this paper, we investigate message compression in the context of large-scale parallel message-passing systems to reduce the communication time of individual messages and to improve the bandwidth of the overall system. We implement and evaluate the cMPImessage-passing library, which quickly compresses messages o ...

**13 QoS provisioning in clusters: an investigation of Router and NIC design**



Ki Hwan Yum, Eun Jung Kim, Chita R. Das

May 2001 **ACM SIGARCH Computer Architecture News , Proceedings of the 28th annual international symposium on Computer architecture ISCA '01**, Volume 29 Issue 2

**Publisher:** ACM Press

Full text available: [pdf\(892.93 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index](#)

terms

*Design of high performance cluster networks (routers) with Quality-of-Service (QoS) guarantees is becoming increasingly important to support a variety of multimedia applications, many of which have real-time constraints. Most commercial routers, which are based on the wormhole-switching paradigm, can deliver high performance, but lack QoS provisioning. In this paper, we present a pipelined wormhole router architecture that can provide high and predictable performance for integrated traffic ...*

**Keywords:** VirtualClock, cluster network, network interface, preemption mechanism, quality-of-service, router architecture, wormhole router

- 14 A study of the impact of direct access I/O on relational database management systems

Heidi Scott, Patrick Martin, Berni Schiefer

September 2002 **Proceedings of the 2002 conference of the Centre for Advanced Studies on Collaborative research**

Publisher: IBM Press

Full text available: [pdf\(117.37 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

*Direct access I/O allows an application program to send requests directly to the I/O subsystem without involving the operating system. We believe that data-intensive applications such as database management systems (DBMSs) stand to reap significant performance benefits from direct access I/O. In this paper we describe an initial attempt to verify this claim. We present a prototype direct access file system and describe a set of experiments we conducted with the prototype and a modified ve ...*

- 15 Parallel architectures: High performance RDMA-based MPI implementation over InfiniBand

Jiuxing Liu, Jiesheng Wu, Sushmitha P. Kini, Pete Wyckoff, Dhabaleswar K. Panda

June 2003 **Proceedings of the 17th annual international conference on Supercomputing**

Publisher: ACM Press

Full text available: [pdf\(222.74 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

*Although InfiniBand Architecture is relatively new in the high performance computing area, it offers many features which help us to improve the performance of communication subsystems. One of these features is Remote Direct Memory Access (RDMA) operations. In this paper, we propose a new design of MPI over InfiniBand which brings the benefit of RDMA to not only large messages, but also small and control messages. We also achieve better scalability by exploiting application communication pattern ...*

**Keywords:** InfiniBand, MPI, cluster computing, high performance computing

- 16 Promises and reality: Server I/O networks past, present, and future

Renato John Recio

August 2003 **Proceedings of the ACM SIGCOMM workshop on Network-I/O convergence: experience, lessons, implications**

Publisher: ACM Press

Full text available: [pdf\(225.62 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

*Enterprise and technical customers place a diverse set of requirements on server I/O networks. In the past, no single network type has been able to satisfy all of these requirements. As a result several fabric types evolved and several interconnects emerged to satisfy a subset of the requirements. Recently several technologies have emerged that enable a single interconnect to be used as more than one fabric type. This paper will*

describe the requirements customers place on server I/O networks; t ...

**Keywords:** 10 GigE, Cluster, Cluster Networks, Gigabit Ethernet, I/O Expansion Network, IOEN, InfiniBand, LAN, PCI, PCI Express, RDMA, RNIC, SAN, Socket Extensions, TOE, iONIC, iSCSI, iSER

**17** [Internet nuggets: Internet nuggets](#) 



Mark Thorsen

March 2003 **ACM SIGARCH Computer Architecture News**, Volume 31 Issue 1

**Publisher:** ACM Press

Full text available:  [pdf\(415.90 KB\)](#) Additional Information: [full citation](#), [index terms](#)

**18** [Simulation and architecture evaluation: Orion: a power-performance simulator for interconnection networks](#) 

Hang-Sheng Wang, Xinping Zhu, Li-Shiuan Peh, Sharad Malik

November 2002 **Proceedings of the 35th annual ACM/IEEE international symposium on Microarchitecture**

**Publisher:** IEEE Computer Society Press

Full text available:   [pdf\(1.14 MB\)](#)  Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)  
[Publisher Site](#)

With the prevalence of server blades and systems-on-a-chip (SoCs), interconnection networks are becoming an important part of the microprocessor landscape. However, there is limited tool support available for their design. While performance simulators have been built that enable performance estimation while varying network parameters, these cover only one metric of interest in modern designs. System power consumption is increasingly becoming equally, if not more important than performance. It is ...

Results 1 - 18 of 18

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.  
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)